checking for line segments at angular intervals of 15°.

In the tree-structured pixel-indexing scheme, the smallest hexagonal pixels in nearest-neighbor groups of seven constitute larger pixels, which, in turn, are grouped into still larger pixels (see figure), and so forth, proceeding from the finest resolution to the coarsest. Each pixel is identified by a sequence of integers in order of increasing resolution. In other words, the first integer in an address denotes the coarsest pixel

that contains the location in question, the second integer denotes whichever one of the seven subpixels of the coarsest pixel contains the location of interest, and so forth down to the last integer, which denotes the smallest hexagonal cell that contains the location of interest.

For some purposes — especially performing quick searches for images that match query images, it is useful to re-sample a binary image data from finer hexagonal grids onto coarser hexagonal

grids. In such a case, each pixel at the next coarser level of resolution is made exactly hexagonal and considered dark if more than three of its seven component pixels are dark.

Finding Minimum-Power Broadcast Trees for Wireless Networks
Algorithms for identifying viable trees have been derived.

NASA's Jet Propulsion Laboratory, Pasadena, California

Some algorithms have been devised for use in a method of constructing tree graphs that represent connections among the nodes of a wireless communication network. These algorithms provide for determining the viability of any given candidate connection tree and for generating an initial set of viable trees that can be used in any of a variety of search algorithms (e.g., a genetic algorithm) to find a tree that enables the network to broadcast from a source node to all other nodes while consuming the minimum amount of total power. The method yields solutions better than those of a prior algorithm known as the broadcast incremental power algorithm, albeit at a slightly greater computational cost.

It is not possible to give more than a highly abbreviated and oversimplified summary of the method within the space available for this article. However, to give meaning to even this brief summary, it is necessary to present some details of the underlying rules, simplifying assumptions, and mathematical constructs in the following two paragraphs.

Each node is equipped with an omnidirectional antenna. The minimum transmitter power that enables the jth node to send information to the ith node is proportional to r_{ij}^{α} where r_{ij} is the distance from node j to node i and α is a channel-loss exponent that usually lies between 2 and 4, the exact value depending on the nature of the signal-propagation medium. Any node in a network can be used to relay a signal to any other node. For the purposes of this method, only the transmitter power levels are calculated: the

receivers are assumed to draw negligible power.

For a network of N nodes, one constructs a power or cost matrix, which is an N-by-N matrix **P**, of which each element P_{ij} is proportional to the minimum power needed for communication between nodes i and j as described above. To represent connections between nodes of the network, one uses an N-element vector, denoted a cut vector, each element of which represents the location of an element in a row of the power matrix. Associated with each cut vector is a threshold vector, the elements of which are the power-matrix elements picked out by the cut vector. The cost of a cut is defined as the sum of the elements of the threshold vector. A cut is regarded as viable if it enables a broadcast to reach all nodes; otherwise, it is regarded as not viable. Another means of characterizing the network is the transfer matrix, which is an N-by-N matrix that is a function of the threshold vector.

One element of the present method is a construct denoted the viability lemma, which provides a necessary and sufficient condition for a cut to be viable. The viability lemma is implemented by equations that require power operations on the transfer matrix. In the case of a large network, power operations on a matrix can be computationally expensive. An alternative algorithm for determining viability does not require matrix power operations; instead, it requires an iterative procedure that involves examination of the rows of the transfer matrix and that is guaranteed to reach an indication of

either viability or non-viability in $N-1$ or fewer iterations.

Another element of the present method is the connection matrix, which is a binary N-by-N matrix representation of a tree. The connection matrix is built iteratively from a cut that has been found to be viable by one of the means described above.

Yet another element of the present method is a stochastic tree-generation algorithm that can be used to generate a viable cut vector. This is an iterative algorithm that starts with a transmission from the source node to a randomly chosen destination node, followed by heuristic selection of the next destination node at each subsequent iteration, until all intended destination nodes are reached. This algorithm converges in $N-1$ or fewer iterations.